

PATENT

-1-

SYSTEM AND METHOD FOR SOURCE IP ANTI-SPOOFING SECURITY

Inventor: Philip Kwan, San Jose, CA

RELATED APPLICATIONS

The present application claims benefit from U.S. Provisional Patent Application Serial No. 60/472,158, filed May 21, 2003, which is incorporated herein by reference. The present application also claims benefit from U.S. Provisional Patent Application Serial No. 60/472,170, filed May 21, 2003, which is incorporated herein by reference.

FIELD OF THE INVENTION

The present invention relates to a method of providing for enhanced security on a computer network to reduce the risk created by the spoofing of IP addresses.

BACKGROUND

As is widely known source IP Address spoofing is a common technique used in denial of service attacks (DoS). Other types of source IP address spoofing attacks are widely known, and include attacks such as distributed denial of service attacks (DDoS), Worm attacks, and Man In the Middle attacks. Spoofed Source IP Address attacks can also include Smurf attacks, NameServer attacks, and ICMP, IGMP, and UDP protocol attacks. One goal in some spoofing attacks is to spread a software virus to as many random new victims as possible, and other attacks are designed to overwhelm a computer system, and other attacks are used to steal information.

Figure 1 shows a computer network 100 of the prior art. At the lowest layer (layer 1 of the OSI Networking protocols) is the physical layer which describes the actual physical elements such as cables and connectors which connect different devices of the computer network. The next layer of the system is the layer 2, the datalink layer. At this level, among other things, the MAC addresses are used to identify the devices which are interconnected on a subnet. As is widely known a MAC address is a unique address which corresponds to a

device connected to a network. The MAC address is generally determined by the Ethernet board of a device which is connected to the network.

The computer network 100 can have a number of subnets. As shown in Fig. 1, the subnets are 102, 104, 106, and 108. Each subnet can contain a number of layer 2 devices such as switches. For example, subnet 102 is shown as having switches 110-126, and subnet 104 is shown as having switches 128-134. The layer 2 devices are not shown for subnets 106 and 108, but as one of skill in the art will appreciate, most subnets will include a number of layer 2 devices such as switches or hubs. Each switch can have a number of ports to which additional switches can be coupled, or to which host devices such as end user computers, or servers, or mainframes can be connected. At the subnet level different devices are connected to the subnet and can communicate with other devices on the subnet by transmitting data packets through the switches. These data packets include the MAC address for the device to which the data packet is to be sent (the destination MAC address) and the MAC address for the host which is sending the packet (the source MAC address). In addition each host device will be assigned an IP address. The IP address is utilized by a router 136 to determine routing for data packets which are being sent by a host on one subnet to a host on a different subnet, or to a different device which may require that the packet be transmitted via the Internet 138. The IP address is often assigned using the Dynamic Host Configuration Protocol (DHCP). Each host on a subnet will normally be assigned an IP address. As is widely known data packets generated by a host on the subnet can include information which is being sent from one host to another host, and further these data packets will include MAC addresses as described above, and a source IP address and a destination IP address.

Source IP spoofing occurs when an attacker host uses a source IP address, which does not correspond, or is not assigned, to its MAC address, in a transmitted data packet. For example, the attacker host may select a source IP address for a different host on a different subnet and transmit a data packet with this false, or spoofed, IP address. This data packet would then be received by the destination device, and the destination device would read the spoofed IP address and it would appear to the destination device that the data packet had come from the device which is actually assigned the source IP address which was used by the attacker host.

In terms of network security defenses, traditional blocks to this type for source IP spoofing were to create inbound filters on the router ports 140-146 that supported the subnets 102 - 108. The router filter operates such that it knows which IP addresses should be received from a specific subnet connected to the particular port. This allows ISP's and enterprises to block randomly spoofed source IP addresses, where the spoofed IP address received on a particular port of the router, is not consistent with source IP addresses for the subnet which is coupled to the particular port of the router. Hackers have recognized the limitations inherent in this type of source IP address anti-spoofing process, and developed spoofing software tools, some of which are referred to as "zombies, and "bots" which now spoof source IP addresses from within their own subnet and subnet mask settings. For customers with large class B subnets, the router level (layer 3) type of defense is not very effective as hundreds and potentially thousands of hosts on the subnet can still be affected.

An Automatic Spoof detector (referred to as "Spoofwatch") has been developed in an attempt to efficiently detect hosts performing source IP spoofing. Spoofwatch works on the premise that these hosts do not respond to ARP requests for their spoofed IP addresses. This solution has many potential shortcomings. For example, the router 136 can receive very large numbers of different source IP addresses in different data packets. Thus, a very large amount of router's processing power is consumed with generating the ARP requests and monitoring the responses.

A review of a number of different websites related to networking showed a number of different approaches related to preventing IP address spoofing, but each approach was very different than that discussed herein. Other techniques have been developed for providing defenses against source IP address spoofing. One of these other approaches relies on using encryption, and source IP filtering at the layer 3 level, which is after the data packets have been transmitted from the subnet to the router.

Additional information regarding different approaches to combating source IP spoofing can be found at

http://www.cisco.com/en/US/tech/tk86/tk803/technologies_tech_note09186a00800a7828.shtml. Additional websites provide discussion regarding the risks associated with source IP

address spoofing and provide some discussion for ways to combat spoofing, see for example: <http://www.sans.org/rr/threats/spoofed.php>; http://www.cert.org/incident_notes/IN-2000-

04.html; http://www.anml.iu.edu/PDF/Automatic_Spoof_Detector.pdf; and
<http://www.linuxgazette.com/issue63/sharma.html>.

In order to increase the efficiency and effectiveness of combating source IP spoofing it would be beneficial to provide source IP spoofing at lower level in the network hierarchy, in a manner which has not previously been provided.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 shows an overview of a system of the prior art.

Fig. 2 shows an embodiment of a network device of the present invention.

Fig. 3 shows a method of an embodiment of the present invention.

Fig. 4 shows a method of an embodiment of the present invention.

DETAILED DESCRIPTION

One approach to improving defenses against Source IP Spoofing is to attack the problem at the subnet or Layer 2 level. Because ISP's and Universities have been hard hit with spoofing attacks, a feature that stops Source IP Spoofing at the Layer 2 subnet level provides a number of advantages. ISPs and Universities frequently have very large subnets, and as a result, utilizing defenses against spoofing at the router level can consume an inordinate amount of the router's processing power. In the past networking devices, such as switches, at the subnet level did not analyze source IP addresses in data packets sent by hosts on the subnet, and in particular it is believed that networking devices at the layer 2 level did not analyze source IP address information in data packets to provide anti-spoofing security procedures based on an analysis of source IP address information in data packets transmitted by hosts on the subnet. Some layer 2 switching devices did provide for some security on ports of the switch, where the source host MAC address was used on an inbound filter on the port to which the host was connected. However, this MAC address type of port security did not provide effective protection against a host attacker that was spoofing source IP addresses.

Fig. 2 shows a view of an embodiment of a network device 200 of the present invention. This network device provides layer 2 switching functions. The network device includes a number ports 202-232. End user host devices (not shown) such as personal computers can be coupled to these ports 202-232, and it is possible for other network devices

such as hubs or additional switches to be connected to a port of the network device 200. A subnet could include one of the network devices 200, or could include a large number of network devices 200 coupled together and connected with a large number of hosts to form a large subnet. The operation of the network device 200 allows for passing data packets received on a port through the network device switching 234 and then transmitting the received data packets through a different port, such that the data packet is transmitted to an intended destination device.

The basic switching operation of such a network device is well known. The network device 200 contains a processing device which operates to analyze data packets received on a port to identify the MAC address of the host sending the data packet. Each data packet can include at least the MAC address of the device sending the data packet (source MAC address) and the MAC address of the device to which the data packet is to be sent to (destination MAC address). As discussed above each host on the subnet can also have an IP address. In the past where one host device on a subnet is sending data packets to other host devices on the same subnet, a switching device would refer to a MAC address look up table to determine which port the destination host was on, and the data packet would be transmitted through the port which is connected with the destination host. A typical layer 2 switch would not analyze the source IP address in connection with this switching function, and would not use the source IP address to provide source IP anti-spoofing operations.

In the network device 200, however, additional functions are provided which allow for utilizing source IP address information. As described herein much of this additional functionality is described in connection with a port security processor 242. Functions of the port security processor can be implemented in a single processor, which is programmed to provide a number of different functions, or aspects of the functions of the port security could be implemented by different processors which work cooperatively to provide the functions herein. As shown in the embodiment of Fig. 2 the port security processor 242 includes a MAC address detector 238 which detects when a new host has been coupled to one of the ports 202-232 of the network device 200; the MAC address for the new host is stored in a table which correlates the MAC address for the new host with the port to which it is coupled. In one embodiment, this MAC address table would be stored in an ACL-CAM discussed in more detail below. In addition when a new MAC address is identified, a source IP address

detector 236 operates to identify the source IP address which corresponds to the MAC address for the new host. This source IP address and the corresponding MAC address are then stored in the table such that each MAC address and source IP address is correlated with each other as a source IP address/MAC address pair.

5 The table which stores the IP address/MAC address pairs can be implemented using a number of different devices. In the embodiment shown in Fig. 2 the table is embodied as an access control list, which are data fields, included in a content addressable memory 240, which is referred to as an ACL-CAM. By utilizing a content addressable memory where the functionality of the memory is determined by hard wiring (as opposed to a CPU which
10 requires the loading of software), the switching of the of data packets is done at a very high speed, and once the MAC address has been determined and the source IP address has been learned, a CPU 244 of the network device 200 is able to operate to monitor and control other aspects of the of the operation of the network device, and the ACL-CAM will control access and switching through the ports.

15 Some prior switches allowed for initially learning source IP addresses for MAC addresses, however, these prior switches were not used for protecting against source IP address spoofing. One limitation of utilizing a prior switch was that these were not designed to allow for the fact that source IP addresses are generally not static, so the source IP address for a MAC address can change over time.

20 In the network device 200, the source IP address detector 236 automatically learns the source IP address for each MAC address entering a port of the network device 200. The port security processor 242 of the network device 200 also provides for dynamically adjusting inbound source IP address anti-spoofing blocking criteria for each port, and a system administrator can specify how many devices or IP addresses to permit per port of the network
25 device 200. For example, the port security processor 242, can be programmed to receive input from a system administrator's computer 246, which can be coupled to the network device 200 by a secure port 248, and to provide information to the system administrator's computer 246. Using the ability to input commands to the network device 200 a system administrator can control aspects of the port security operation, as well as other aspects of the
30 operation of the network device 200. For example, the system administrator could control the maximum number of source IP addresses which are learned from a port, and the network

device 200 would reject any data packets with new source IP addresses that exceed the maximum number.

In one embodiment, the port security processor 242 will periodically poll ports for the learned IP addresses which are stored in the table to ensure that the host devices with the learned source IP addresses are still coupled to the port. If it is determined that a host device having the learned source IP address is no longer coupled to a port then source IP address for the host that is no longer present can be removed from the table so as to allow a new source IP address to have access on the port.

The network device 200 extends port security features beyond the MAC address filtering procedures that were used in prior layer 2 devices. The port security processor 242 allows source IP anti-spoofing protection to be activated selectively on a port-by-port basis. The port security processor 242 uses the source IP address detector 236 to automatically learn the source IP addresses for each host device attached to the port. To determine if the data packets received at a port contain a new source IP address that has not been learned, the ACL-CAM compares the source IP address and the MAC address in a received data packet, with the table of IP address/ MAC address pairs.

When a new MAC Source Address is detected on a port by virtue of a received data packet identifying a new MAC source address, the source IP address detector 236 learns the association of source MAC address and the corresponding source IP address. Once the pair is learned, the ACL-CAM 240 is programmed with the information and switching of the network device proceeds to switch data packets normally.

The source IP address detector 236, can learn the source IP address for a host in a number of different ways. For example, the source IP address can be learned by using a reverse address resolution protocol (RARP) which provides for sending out the MAC address on the subnet, and in response the host having the MAC address responds with its IP address. The source IP address detector 236 could also learn the IP address for a host by listening to the DHCP response packet being returned to the host. This response contains the source IP Address for the host. When a DHCP packet is detected, the entry in the table for the MAC address receiving the DHCP packet with a source IP address is cleared and the source IP address provided in the DHCP packet is loaded into the table. This utilization of the DHCP response works well where the port is set to allow for one IP address. It should be noted that

if a port is set to allow for more than one IP address for the port, then relying on the DHCP response alone may be insufficient, as the DHCP response may not allow for an unambiguous correlation of the source IP address to the correct MAC address. However, this ambiguity could be resolved if the DHCP request and the DHCP response were both tracked.

5 Another technique provides for watching for the IP header information in a data packet when the host first transmits a data packet through a port. If a static IP address is used and the port is set to 1 IP Address, the user can unplug the cable (causing a link down) to reset the table, or a timer can be used to clear the table. Another technique provides for trapping (listening to) ARP requests and ARP reply messages to learn the source IP address and MAC address
10 pairs, and storing the pairs in ACL-CAM for each port.

In addition to the above, it should be recognized that care must be taken to learn the correct source IP address. For example, in order to support delayed technologies, such as the widely known IEEE 802.1X standard which may postpone IP Address assignment until the port is authenticated, a tiered, or delayed, approach is required to successfully detect the
15 source IP address. This approach could allow for first identifying the MAC address and then waiting for an appropriate amount of time to learn the correct IP address. It is important to avoid learning the IP addresses assigned by Microsoft or Apple operating systems which may be provided to a host when a DHCP server cannot initially be found. It can be advantageous to delay IP address learning process until a certain amount of traffic has passed through the
20 port. Further the risk of possibly learning the wrong IP address can be reduced by allowing an system administrator to seed the IP learning process with the IP address properties that are consistent with the subnet to which the host is coupled. For example, if the IP Subnet is 10.32.1.0/24, then this information can be used to seed the IP learning process and only match a MAC address with an IP addresses consistent with the subnet 10.32.1.0/24
25 addresses. This will eliminate false learning of a default 169.0.0.0 addresses assigned by Microsoft or Apple operating systems when a DHCP server is not initially located. It is also advantageous to confirm that the learned IP address is correct by performing a reverse IP check.

Once the learned IP address has been confirmed, and the IP address/MAC address
30 pair has been stored in the ACL-CAM, blocking procedures are applied to the port. If there is more than one host device allowed per port, this process of determining MAC address and

source IP address pairs will be repeated for each learned IP address on the port. Once a new source IP Address is confirmed, the inbound blocking procedures are applied to the port and includes the new source IP Address. Additional MAC address and IP address pairs will continue to be learned until the maximum number of MAC and IP Addresses is reached.

- 5 The source IP anti-spoofing procedures should be compatible with existing MAC address port security and 802.1X Port Authentication features. In order to provide for compatibility the following order of execution can be used. First, MAC address port security is utilized, which confirms that a MAC address present on a port is a valid MAC address, and if it is not MAC port security procedures block data packets from the invalid MAC address. If 802.1X
- 10 Port Authentication is enabled, the user will be prompted for the 802.1X user credentials to authenticate the port and either permit or deny data packet traffic based on the success or failure of the IEEE 802.1x authentication process. Source IP Security is then used, if enabled, to check if the maximum number of source IP addresses has been learned for a port. If the maximum number of source IP address has not been learned for the port, then the
- 15 source IP address will be learned and confirmed, and the pairing of the source IP address with MAC address will be set in a table.

In order for the blocking to be efficient and fast, it should be implemented in hardware, such as a CAM having an ACL, as opposed to using a CPU where the operation would be slower. The source IP anti-spoofing methods can provide for different types of

20 security. For example, one aspect of the operation described herein provides for allowing an system administrator to set a maximum number of source IP addresses for a port. By limiting the number of the source IP addresses which can transmit data packets through a port of the network device 200, the risk of certain types of spoofing attacks (such as DoS attacks) can be prevented. This aspect of the operation provides for blocking of data packets

25 at the port of the network device based on the source IP address contained in the data packet. For example, by limiting the number of source IP addresses on a port, an attacking host could only use a limited number of spoofed IP addresses before the maximum number of source IP addresses for the port would be exceeded.

A second operation of the network device 200 blocks data packets at the port, where

30 the data packet contains a source MAC address and source IP address pair which does not match one of the previously identified MAC address/ IP address pairs stored in the table. In

this operation, a host device which is attempting to use a source IP address which does not match the correct source IP address for the MAC address will be blocked at the port. This second part of the operation provides for a high level of security against source IP spoofing attacks.

5 It should be noted that embodiments of the methods and systems herein can be provide for a significant amount of flexibility, which can provide a system administrators and ISPs with a powerful tool to combat source IP spoofing. For example, as discussed above, in one embodiment the number of source IP Addresses which can be associated with each switch port can be selected by an administrator. If more than one source IP address is
10 permitted per port, then source IP spoofing is possible if the attacking host is using a validated MAC address, unless the source IP Security process as been activated to provide for port access based on correlating the learned source IP Address to its MAC address, such that access on a port is blocked, or permitted, based on the matching of the MAC/IP address pair in a received data packet with a MAC/IP address pair stored in a table. It should be
15 recognized, however, that even without providing security based on the MAC/IP address pairs, some degree of protection against Source IP spoofing is provided by allowing the number of source IP addresses on the port to be controlled.

 The port security processor 242 can also include a source IP age out timer 250. The port security processor 242 can allow an administrator to specify a time period for a source
20 IP age-out timer. This timer can clear the ACL-CAM, or other possible table, of source IP addresses, every n seconds to allow the network device 200 to support downstream hubs and switches for multi-host configurations. If a timer were not provided then source IP addresses which were previously on the port might prevent new source IP addresses from gaining access to the port, where the maximum number of source IP addresses would be exceeded.
25 In short, it would not be practical to maintain a link for a source IP address indefinitely. Where an administrator did not want to have a time out entry they could specify a zero "0" for the timer, and a default of 60 or 120 seconds or other appropriate time could be provided. Additionally, or alternatively, the source IP age out timer could also be flow based, which would provide that if the flow stops for a source IP address, for a period of n seconds, then
30 the source IP address can be removed (aged out) from the table.

An embodiment of the system and method can also provide for capturing the information when a possible IP spoofing attempt has occurred. This information could then be used to generate syslog messages which could be transmitted and recorded in a log to record information regarding the operation of the system, including possible IP spoofing attempts.

As shown by the discussion above, it is important in many applications of the source IP security operations herein, that the table containing the learned source IP addresses be dynamic, such that the table can be updated so that timed out source IP addresses can be removed and new learned source IP addresses can be added to the table. Further, being able to change entries in the table allows for the fact that the source IP addresses assigned for a given host can change over time. This means that when the learned source IP addresses are stored in an ACL-CAM, these addresses should not be saved when a write only memory operation is performed.

The port security processor 242 should be programmed such that an administrator can view the source IP addresses learned and/or assigned to the port. Further, the IP security device should allow an administrator to view the setup and configuration of the timers and source IP address ACL's. The port security processor 242 should provide commands which allow an administrator to clear a single entry from the table, and to allow the administrator to clear all entries from the table. Further, debugging tools may also be provided to allow administrators to troubleshoot the security procedures for their particular environments.

Fig. 3 is a flowchart showing a method 300 of an embodiment of the method herein. The method provides for receiving 302 a data packet from a host on a port of the network device. The data packet is analyzed 304 to determine its source IP address. The determined source IP address is then compared 306 with the source IP address, or addresses, which have previously been learned and stored in a table as source IP address which is permitted access to the network through the port. If it is a previously learned IP address then the data packet is passed 308 through the port. If it is not a previously learned source IP address, then it is treated as a new source IP address. Once it is determined a new IP address is on the port, then it must be determined 310 if the maximum number of source IP addresses are present on the port. If the maximum number of source IP addresses are present on the port, then the data packet with the new source IP address is blocked 312 and the information relating to the

blocking of the port is transmitted in a trap and syslog message. If the maximum number of the source IP addresses are not present on the port, then a learn source IP routine is performed 314 (various methods related to this are described above). Typically the learn source IP routine will include doing a reverse IP check to confirm the source IP address. If the confirmation or learning of the source IP address fails then the data packet containing the new source IP address is blocked 316 and syslog message can be generated. If the learning and confirmation of the new source IP address is successful, then the new source IP address is stored 318 in a table indicating that the new source IP address is permitted access through the port. As discussed above this storing of the new source IP address in a table, can include storing the information in an ACL-CAM. Further a syslog message could also be transmitted in connection this operation. At this point the data packet with new source IP packet can be passed 320 through the port. Further, an additional operation not shown in the flow chart includes the operation of the age-out timer which would provide for removing previously learned source IP addresses which are stored in the table and determined to no longer be present on the port.

Fig. 4 shows a method 400 of another embodiment of the invention. The method 400 provides for receiving a data packet on a port and determining its source MAC address 402. The determined source MAC address is compared 404 with MAC addresses shown in a table 404. A determination 406 is made as to whether the source MAC address is new, which would mean that it is not in the table. If the source MAC address was previously stored in the table, then the source IP address and MAC address pair for the received data packet is compared 408 with the source IP address/MAC address pairs in the table. If the pair for the received data packet is found in the table then the received data packet is passed 410 through the port. If the pair for the received data packet is not found in the table, then the received data packet is blocked, or dropped 412 at the port.

If the received data packet at the port has a MAC address which is new, then the source IP address for the received data packet is learned 414 using one of the processes described above. After the source IP address has been learned, a reverse IP check 416 is done to confirm the source IP address. If the reverse IP check is successful 418, then the table is programmed 420 with the IP address/MAC address pair, and the packet is passed

422. If the reverse IP check is not successful then the received data packet is blocked 424, or dropped at the port.

While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example, and not limitation. It will be apparent to persons skilled in the relevant art that various changes in form and detail may be made therein without departing from the spirit and scope of the invention. This is especially true in light of technology and terms within the relevant art(s) that may be later developed. Thus, the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.